Large Scale Novel Object Discovery in 3D

Siddharth Srivastava, Gaurav Sharma and Brejesh Lall

arXiv:1701.07046v1 [cs.CV] 22 Jan 2017

***Abstract*—We present a method for discovering objects in 3D point clouds from sensors like Microsoft Kinect. We utilize supervoxels generated directly from the point cloud data and design a Siamese network building on a recently proposed 3D convolutional neural network architecture. At training, we assume the availability of the some known objects—these are used to train a non-linear embedding of supervoxels using the Siamese network, by optimizing the criteria that supervoxels which fall on the same object should be closer than those which fall on different objects, in the embedding space. We do not assume the objects during test to be known, and perform clustering, in the embedding space learned, of supervoxels to effectively perform novel object discovery. We validate the method with quantitative results showing that it can discover numerous unseen objects while being trained on only a few dense 3D models. We also show convincing qualitative results of object discovery in point cloud data when the test objects, either specific instances or even their categories, were never seen during training.**

I. Introduction

Object discovery is motivated by the need for enabling robots to have human-like capabilities. Since humans perceive information from multiple sensory inputs, contemporary research aims at utilizing cues from different modalities as well [1], [2]. Recent studies [3] show that a majority of these sensory inputs work in conjunction with human visual system —vision consumes most of the brain's processing power. A very important modality for visual perception in humans is visual data in 3D (cf. 2D images) [16]. Such data is relevant to a larger spectrum of multimedia applications which could be broadly classified as (i) those extracting meaningful interpretations from 3D data or (ii) those involving manipulation of 3D data itself. The former involves object recognition [4], [5], 3D Object Retrieval [6], prediction of new views using existing 3D models [7], gesture recognition [8], action recognition [9], [10], scene recognition [11] etc., while the latter includes compression [12], [13], object modelling [14], encoding for efficient processing [15] etc. The problem addressed in this paper falls in the former class; we propose a learning based approach to interpret and understand the meaning implied by the 3D data. The method deals with large scale of data, e.g. the number of points processed in the 3D datasets are $\emptyset(10^6)$, learns models which have a large number of parameters $\emptyset(10^5)$ and discovers objects from approximately 300 unseen object categories—to process such large scale 3D data, efficiency is a key concern, and we demonstrate that the proposed method is able to cope with such scale of both data and models.

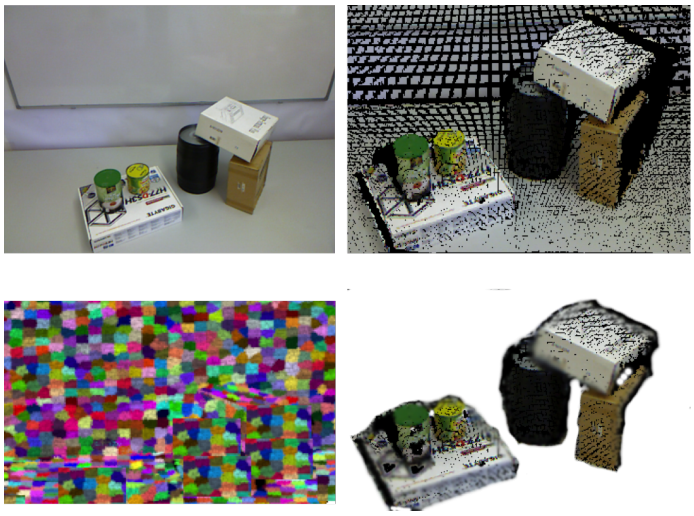

Fig. 1: (Left-top) The RGB image of a scene containing multiple objects. (Right-top) The point cloud representation of the same scene (noisy). (Left-bottom) Supervoxels obtained from the scene. (Right-bottom) Discovered objects (the point cloud has been post-processed for better illustration)

As humans perceive the world in 3D [16], the ability to reason in 3D becomes a necessary computer vision capability. For instance, in case of autonomous mobile robots, an important problem is the efficient navigation of robot in the real world scenario–where inferring and demarcating manipulable objects and their context is one of the main technologies required. Traditionally, such object estimation problem has mostly been well studied in a supervised *known object* scenario [17], [18], i.e. the robot is expected to encounter the same objects it was trained on. A more realistic and challenging setting is when the objects that the robot will encounter were never seen before. In the present paper, we address such important problem of discovering never before seen objects from 3D data. This problem which is usually termed as *Object Discovery* is a hard, relevant and important problem in the domain of perception using data in 3D modality.

Object discovery has been a popular research area in RGB images [19] and has recently been explored for applicability on 3D data [20], [21]. Karpathy et al. [20] recently proposed a method to address such problem in 3D meshes based data, while Mueller and Birk [21] addressed the problem with 3D point cloud data, from sensors such as the very popular Microsoft Kinect. Majority of the methods dealing with Object Discovery in 3D scenes work on the premise that labelled 3D data is scarce and hence look towards unsupervised techniques

Siddharth Srivastava is with the Department of Electrical Engineering at the Indian Institute of Technology Delhi. *eez127506@iitd.ac.in*

Gaurav Sharma is with the Department of Computer Science and Engineering at the Indian Institute of Technology Kanpur. *grv@cse.iitk.ac.in*

Brejesh Lall is with the Department of Electrical Engineering at the Indian Institute of Technology Delhi. *brejesh@ee.iitd.ac.in*

or use very small amounts of training data. With the increased popularity of cheap and commercial RGB-D sensors like Microsoft Kinect, the availability of real world data in 3D is rapidly increasing [22]. Additionally, high quality CAD models of various objects, such as those in ModelNet dataset [23], now provide dense 3D models of various common and important objects. Hence, we propose to directly learn from the available RGB-D images and high quality CAD models to discover unseen objects in RGB-D images, which can be acquired independently at test time.

The ability to discover unknown objects, will aid autonomous robot navigation, especially in new environments, albeit with common objects. While, informed by the traditional wisdom, one would expect that this would require fitting some kind of prior models to the full 3D scenes, as well as the object present therein [5], we instead promote a learning based approach. We propose bottom up aggregation of basic supervoxels, in the 3D scenes, into objects. We still work in the setting where at test time, on the field, the robot would encounter novel objects which were never seen while prior learning. The novel perspective comes from the use of discriminative metric learning approach for obtaining the similarity between supervoxels which are expected to belong to the same object, which might itself be unknown. Recently, Karpathy et al. [20] investigated along similar lines, where they argued that simple properties would distinguish objects from context, e.g. 'a mug on a table can be identified . . . based solely on the fact that it is a roughly convex, symmetrical shape sticking out from a surface'. Here, instead of hard-coding such intuitions, we take the automatic data-driven route and aim at learning such aspects (of supervoxels) which might make them similar to one another and taken together would form an object. We assume that such aspects, while learned on a given set of training objects, would generalize to novel unknown objects and find this to hold reasonably well in our extensive empirical evaluations.

Related to our interests, there have also been recent attempts to find 3D bounding boxes around objects, e.g. generic object proposals in 3D [24], [25]. Here, we aim at discovering strict object boundaries in 3D, and not just bounding boxes. This is more helpful for real world scenarios where scenes are usually cluttered and contain many unknown objects of various sizes, and the robot wishing to manipulate such objects would require finer 3D understanding of the object to e.g. estimate better grasping configurations.

Recently, deep learning based methods have been effectively able to exploit the available annotated data and perform well in different tasks such as object recognition, shape completion etc. [26], [27], [23], [4], [28]. In line with those successes, we propose to use a deep Siamese network for the task of learning non-linear embeddings of supervoxels into a Euclidean space, such that the ℓ_2 distances in the embedding space reflects the object based similarities between the voxels. With the learnt embeddings, we enforce that the supervoxels which originated from the same object (category) come close to each other and those that originated from different objects, move far apart. Such non-linear (metric learning via) embeddings has been recently very successfully applied to various problems of computer vision, e.g. face verification [29], face retrieval [30] and unsupervised feature learning [31], [32], we explore ideas informed by these methods in the domain of 3D computer vision, specifically working with popular point clouds modality. This work thus builds on the very recent work using deep networks for 3D data, e.g. VoxNet [27], and uses them for the more challenging setting of novel object discovery. To the best of our knowledge, we are the first ones to show that a deep architecture can be successfully designed for this task, using fully volumetric data and trained end-to-end with good performances in challenging benchmarks, outperforming existing approaches.

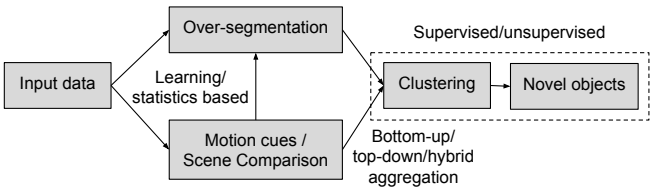

Fig. 2: General Pipeline of Object Discovery Methods

In summary, we address the problem of 3D object discovery in a supervised, but restricted setting where we are given a few known objects at train time, while, at test time we are expected to discover novel unknown objects. The contributions of this paper are as follows. (i) We propose a Siamese network based embedding learning framework for supervoxels which works directly with point cloud data and can be trained end-to-end, (ii) the proposed method is fast, requiring three sub-steps of (a) supervoxel computation, (b) a CNN forward pass and (c) classification of supervoxels with learned discriminative embeddings, (iii) we show with quantitative and qualitative results, on challenging public datasets of 3D point clouds, that the structure of objects in terms of distances between supervoxels can be learnt from only a few reference models and can be used to generalize to (supervoxels of) unseen objects, (iv) we compare with many relevant existing approaches for the task on public benchmarks with large number of classes and object instances per class (~ 100 and ~ 200 respectively) and report state-of-the-art results.

II. Related Work

We now discuss some of the works which are closely related to the problem addressed and/or to the proposed method.

A. Object Discovery in 3D

The general pipeline followed by the majority of the works related to object discovery in 3D is shown in Fig. 2. Among the representative works, Karpathy et al. [20] discover objects by using multiple views from a 3D mesh. They over-segment the meshes and obtain objectness scores to rank various segments. Garcia et al. [33] generate object candidates by over-segmenting the RGB-D images. They then use Voxel Cloud Connectivity Segmentation (VCCS) [34] for over-segmenting the RGB-D image into voxels. They then fuse the segments with various ranking criterion, to obtain discovered objects.

Compared to these works, we apply discriminative metric learning to learn which segments should be fused. Meuller et al. [21] first over-segment the RGB-D point cloud and then learn a classifier for primitive objects. The complex objects are formed as composition of primitive objects. The approach is similar to ours but in our method, we do not explicitly assume any structural composition of basic shapes that constitute an object, allowing us to deal with more complex and novel object categories. Richtsfeld et al. [35] perform object discovery on the constrained domain of table-top scenes. They pre-segment a scene based on surface normals and subsequently construct surface patches followed by training an SVM for learning the relation between patches on same object. Firman et al. [36] learn similarity among segments of an image by training a classifier as well. They cluster the segments using a variation of correlation clustering. They use handcrafted features for segments to eliminate non-discriminative information. In contrast, our features are learned using a deep convolutional network trained on CAD models and we learn the similarity and dissimilarity among segments (supervoxels) with a supervised metric learning based objective.

In a different direction cf. the above works, Herbst et al. [37] construct 3D maps from multiple views corresponding to the movement of a robot in an environment. They detect objects that move among scenes and use an iterative algorithm with scene differencing to match objects among scenes. In another work [38], they propose to discover objects by determining frequently occurring object instances and determining the patches that move in the scenes. In comparison, the essence of our method lies in characterizing properties of objects using CAD models. Shin et al. [39] find repetitive objects in 3D point clouds by forming superpixel like segments of points based on relation between normals of neighbouring points. They define a distance metric among segments as the distance between a weighted combination of various descriptors. This is followed by a joint compatibility test to determine segment associations. As a limitation, they can only find objects which are repeated in a dataset as their method is based on finding multiple occurring segments. We overcome this limitation by learning the discrimination measure among supervoxels belonging to same and distinct objects.

Bao et al. [40] perform attention driven segmentation by first detecting a salient object using depth information and then forming a 3D object hypothesis based on the shape information. They label and refine objects using Markov Random Field in voxel space with constraints of spatial connectivity and correlation between color histograms. Stein et al. [41] decompose a scene into patches while categorizing edges as convex or concave and thus creating a locally connected subgraph. Their technique is unsupervised but mostly detects object parts instead of complete objects. Gupta et al. [42] perform segmentation based on an object's geometric features such as pose and shape. Asif et al. [43] propose perceptual grouping of segments based on geometric features. While most of the previous methods cater to indoor scenes, Zhang et al. [44] tackle the problem of discovering category in an urban environment. They pose object discovery as a problem of determining category structures by finding repetitive patterns of shape and refining the segmentations by defining an energy function which minimizes distances between the shape patterns.

B. Supervoxels

Our model learns to reason at the level of supervoxels for the task of novel object discovery; we thus briefly review some works on supervoxel computation. In 3D data, supervoxels are groups of voxels having some coherent properties such as boundary, texture, color etc. similar to superpixels for 2D images [45]. Papon et al. [34] present Voxel Cloud Connectivity Segmentation (VCCS). The method involves converting the point cloud into a voxel cloud and building a voxel adjacency graph using a voxel kd-tree, followed by spatial seeding where candidate supervoxels are chosen. For this, the space is partitioned into a voxel grid and the initial seeding candidates are chosen as the voxels having minimum distance to the centre of occupied seeding voxels. This may result in some isolated voxels, which are filtered out by enforcing a voxel occupancy criteria within a small search volume around the seed voxels. The remaining voxels are moved to a connected voxel with the smallest gradient in the search volume. This is followed by initialization of a supervoxel feature vector in a 39 dimensional space with the centre of the seed voxel and two neighbours. The voxels are then clustered by assigning them to nearest supervoxels moving to next voxel in a breadth first search manner in the voxel adjacency graph. The process is repeated till the cluster centres stabilize. In addition to segmentation related tasks, VCCS has found applications in many fields such as Scene Reconstruction [46], Saliency Detection [47], Scene Understanding [48] etc.

The supervoxel clustering discussed above works directly on 3D point clouds. However, there have been extension of 2D superpixel based techniques to 3D volumes [45], [49], [50], [51]. Weikersdorfer et al. [52] propose Depth Adaptive Superpixels which works in 2.5D and performs clustering based on depth and normals. Mueller et al. [21] extend the idea of VCCS to develop over-segmentation based simple and complex objects. The complex patches constitute the information from simpler patches or primitive shapes.

C. 3D Convolutional Neural Networks

Maturana et al. [27] propose a 3D-Convolutional Neural Network (3D-CNNs) architecture, called VoxNet, which can work with volumetric occupancy grid. An advantage of this architecture is that it can work with many types of 3D data such as LiDAR, RGB-D, Voxelized Clouds etc. Although the method is tuned to work for the task of object recognition where the objective is to predict the category label for an input point cloud, in this work, we build upon the architecture (Sections III and IV) to perform object discovery by learning a distance metric between supervoxels. Wu et al. [23] propose 3D Shapenets using Convolutional Deep Belief Networks and demonstrate its effectiveness in shape completion and object recognition. Maturana et al. [53] apply 3D-CNN for the task of landing zone detection from LiDAR images. Garcia et al. [54] develop an architecture to classify objects in real time using

CNN on point cloud data while Huang et al. [55] label point cloud using a 3D-CNN. Qi et al. [56] study volumetric and multi-view representations of CNN for object classification. While there have been works using deep learning for 3D data, in the present work we propose a novel Siamese network based framework to learn highly generalizable distance metric between supervoxels abstracting out object properties, i.e. the metric is able to predict (a small distance) if two supervoxels belong to the same object or not, even for objects which were not seen during training. We then use the learnt metric for the challenging task of novel object discovery in 3D.

In addition to neural networks for classification, recently Siamese networks have also been used for a variety of computer vision tasks, such as face matching [57] and semantic image retrieval [58].

III. Approach

We now describe our approach for discovering objects in 3D point cloud data. Fig. 3 gives the overall block diagram of the approach at test time. At training time the 3D convolutional neural network is trained in a specific manner which we explain below. We obtain point cloud data from an appropriate sensor, e.g. Microsoft Kinect, and then proceed as follows. The overall strategy is to discover novel objects in 3D by clustering supervoxels obtained from the point cloud of a scene. To obtain the distance metric required for clustering, we propose to use deep Siamese network for learning non-linear embeddings of supervoxels into a Euclidean space, where the distance reflects the object membership, i.e. the supervoxels which belong to the same object are closer than those which belong to different objects. Learning such a distance metric requires supervised training data, which we obtain from (the point cloud data of) a set of few known objects. Note that while the proposed method requires annotated training data, the data it uses is not from the objects it will see at test time, i.e. at test time the method is expected to discover novel objects which it had not seen at training. Hence, the method is different from the traditional object detection methods which aim to find the objects, which they were trained on. Finally, when we have the object based supervoxel embedding, we perform clustering using an efficient method to obtain object hypothesis. With this overview of the method, we now explain each component in complete detail.

A. Oversegmentation using Supervoxels

The first step of our method involves obtaining an oversegmentation of the 3D data using supervoxels. Towards that goal, we use an efficient and robust method to deal with the noisy and large amount of point cloud data. We group the points in the point clouds into supervoxels using Voxel Cloud Connectivity Segmentation (VCCS) approach of Papon et al. [34]. Since (i) the number of points in the point clouds are very large, $\emptyset(10^6)$, and (ii) they are noisy as well, this first step of supervoxelization, reduces the computational load on the successive stages and also lends robustness to noise. Since the objective optimized by the supervoxels creation method is that they should not cross object boundaries, we aim to thus cluster these supervoxels to give us plausible novel objects at test time. The seed resolution used to construct supervoxels describes the fineness (low seed resolution) or coarseness (high seed resolution) of the supervoxel segmentation, allowing us to have multi-scale supervoxels that can capture fine details as well as larger blocks of the objects of interest. Hence, we extract supervoxels at multiple seed resolutions and later merge them appropriately (described below) to coherent object hypotheses.

B. Siamese Deep Network

As the next component, we intend to have a learnt similarity metric which captures object based distances between the supervoxels, i.e. it brings those on the same object closer while pushing those on different objects far. To do this, we employ a Siamese network which learns a non-linear embedding into a Euclidean space where the distance captures such semantic relation between the supervoxels.

Now, we first outline how we utilize the supervoxels, generated above, with the Siamese network and then describe the network architecture, loss function as well as training procedure in detail.

Once we have the supervoxels from the point cloud data, we generate 'positive' and 'negative' pairs (for each seed resolution) for training the Siamese network as follows. To generate the positive pairs of supervoxels, we consider all the objects that are labelled during training. For all the supervoxels $\{z_i | i = 1, \ldots, N\}$, in the training set, we create two sets containing all possible pairs of supervoxels of the following two kinds respectively, (i) the two supervoxels lie on the same object (S^+) or (ii) they lie on different objects or background (S^-). To decide if a supervoxel z_i lies in an object $x \in \mathcal{X}$ (with $\mathcal{X}$ being the set of all objects in the current scene) or not, we take the intersection of the points z_i with the set of points lying on the object x. If the fraction of these points is more than a threshold $\beta \in \mathbb{R}$, of the total number of points in the supervoxel, then it is considered to lie on the object. Note that by doing this we are, effectively, not doing a hard assignment where the supervoxel is required to lie completely in the object, but doing a soft assignment where the supervoxel is required to be sufficiently inside the object. Similarly, the set, S_-, of negative pairs contains supervoxels which do not belong to the same object, i.e. they either belong to two different objects or to an object and to the background, respectively. More details on selecting supervoxels on object and background for training the network are provided in Sec. IV. The following equation formalizes this mathematically.

$$
\begin{aligned}
&\forall i, j \in [1, N],\ i \neq j \\
&(z_i, z_j) \in S_+ \quad \text{iff, } \exists x \in \mathcal{X} \text{ s.t.} \\
&\frac{|pts(z_i) \cap pts(x)|}{|pts(z_i)|} \geq \beta \text{ and } \frac{|pts(z_j) \cap pts(x)|}{|pts(z_j)|} \geq \beta, \\
&(z_i, z_j) \in S_- \text{ otherwise}
\end{aligned} \tag{1}
$$

where, $pts(t)$ gives the set of points in the object t and $|\mathcal{A}|$ denotes cardinality of the set $\mathcal{A}$.

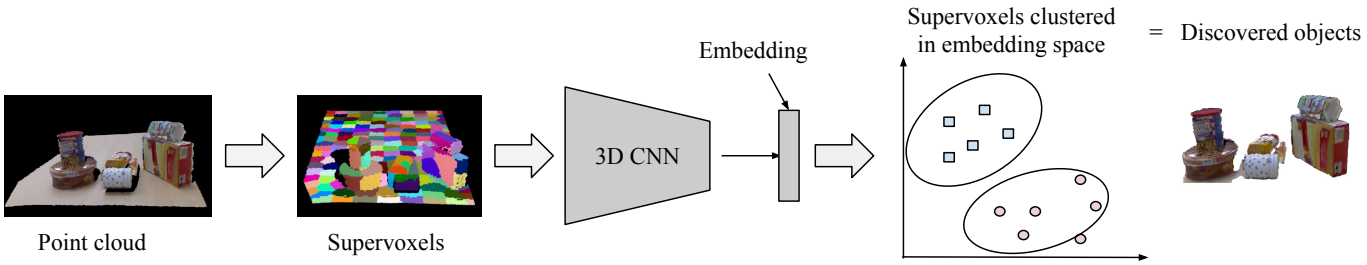

Fig. 3: Block diagram of the proposed approach at test time. The 3D convolutional network used to embed the supervoxels is trained using a Siamese network which captures object based similarity of the supervoxels. See Sec. III for details.

Given such pairs of same and different object supervoxels, we proceed to train the Siamese network. Our Siamese network is based on the recently proposed convolutional neural network for 3D point cloud data by Maturana and Scherer, called VoxNet [27]. VoxNet takes the point cloud as input, computes the occupancy grid with it and transforms it using two convolutional layers, a max pooling layer and finally a fully connected layers to finally perform classification over a predefined number of objects. While it was proposed for object classification in 3D we use it for processing supervoxels for obtaining their semantic non-linear embeddings. We have two parallel streams of the VoxNet, with tied parameters, which take the two members of the positive (or negative) pair and do a forward pass up to the last fully connected layer (removing the final classification layer). On top of the last fully connected layer from the two streams, we put a loss layer based on the hinge loss for pairwise distances given by,

$$\mathcal{L}(\Theta, W) = \sum_{S_+ \cup S_-} \left[b - y_{ij}(m - \|W f_\Theta(z_i) - W f_\Theta(z_j)\|^2 \right]_+$$
$$\text{where, } [a]_+ = \max(a, 0)\ \forall a \in R. \quad (2)$$

The VoxNet forward pass, parametrized by Θ, is denoted by $f_\Theta(\cdot)$, W is the parameter of the projection layer after VoxNet, (z_i, z_j) is the pair of supervoxels and $y_{ij} = +1$ or -1 indicates if they are from the same or different objects, i.e. if the pair belongs to S_+ or S_-, respectively. $b \in \mathbb{R}$ and $m \in \mathbb{R}^+$ are bias and margin hyper-parameters respectively.

The loss effectively ensures that in the embedding space obtained after the forward pass by VoxNet and the projection by W, the distances between positive pairs are less than $b-m$, while those between the negative pairs are greater than $b+m$. When trained on the given train objects, we expect the loss to capture the semantic distances between the supervoxels in general, which would then generalize to supervoxels of novel unknown objects at test time. As we show in the empirical results, we find this to be indeed the case.

C. Supervoxel Clustering and Postprocessing

Once we have trained the Siamese network, we can embed the supervoxels into a object based semantics preserving Euclidean space. We then perform clustering in this space to obtain novel object hypotheses. We use DBSCAN [59] to group similar supervoxels into larger segments. DBSCAN clusters the supervoxels based on density i.e. distribution of nearby supervoxel features. As the features are learned by the deep siamese network to be closer in the embedding space if they belong to same object, they are considered as potential neighbours by DBSCAN and vice-versa for those which do not belong to the same object.

We observed that most of the objects were spanned by multiple supervoxels, especially since the supervoxels were extracted at multiple scales. Therefore, the criteria for a dense region in DBSCAN was specified as at least two supervoxels big. This limits the method to discover objects that consist of at least two supervoxels. In our experiments, we did not find any object with single supervoxel in the training set, but we do note that in the real world where a robot is navigating, a single supervoxel may possibly correspond to a heavily occluded object. But it would be important to note that in case of robot navigation, detecting such heavy occlusions is not a problem since the the object occluding the view needs to be processed first before moving on to the object behind it. Moreover, in such cases there are chances that the visible patch of the occluded object becomes a part of other object. As we show in experiments, the proposed technique is highly robust to such patches and avoids intermixing among object boundaries.

IV. Experimental Results

In this section, we discuss the empirical results of various experiments to analyse the proposed method. We begin with the description of the datasets used followed by implementation details and baseline. Then we proceed towards discussing the results on these datasets.

A. Experimental Setup

Datasets. We demonstrate the effectiveness of the technique on multiple challenging and publicly available datasets. The dataset considered in the evaluations are as follows.

1) *NYU Depth v2* [60] RGB-D dataset consists of 1449 RGB-D images with 894 classes and 35064 object instances. We use this dataset to create a challenging test setting for object discovery. This is contrast to most of the previous works where evaluation was done on much smaller datasets consisting of only a few hundred object instances belonging to only a few classes. Additionally, the scenes in NYU Depth v2 are from complex real world indoor scenarios with high depth variations,

making it more challenging compared to other datasets used for evaluating object discovery methods previously. We select NYU Depth v2 as it consists of complicated real world indoor scenes and has ground-truth pixel level annotations, making it ideal for setting up a challenging benchmarking setup.

2) *ModelNet dataset* [23] was introduced by Wu et al. . It has two versions, i.e. ModelNet40 and ModelNet10. As the names suggest, ModelNet40 contains $151,128$ 3D models and has 40 object classes while ModelNet10 has 10 object classes . We use the standard training set provided with the dataset for pre-training 3D-CNN in our experiments.
3) *Universal Training Dataset.* We create a training set from NYU Depth v2 dataset with 60% of the classes. We use this as the universal training set for all the experiments reported in this paper. In order to capture the essence of object discovery, we ensure that none of the objects in evaluated test datasets have been seen before. We note explicitly here that the *Universal Training Dataset* does not contain any class that are present in any of the test data. During testing we only use test split of the corresponding datasets and report results on them except for NYU Depth v2 where we use the remaining classes (and instances) for computing the results while testing. As a scene in NYU Depth v2 consist of multiple classes and its instances, any object belonging to the test class is labelled as clutter/background in the training set. We first tried to take only images with the train objects and no test class objects but could not create a big enough training set, hence we resorted to the current setting. In out training set while some of the test objects may be present, they are not annotated (marked as background) and are never used for the supervised training.
4) *Object Discovery dataset* [61] was proposed in 2016 specifically for tackling the problem of object discovery. The dataset consists of 30 RGB-D scenes with 640x480 resolution. There are total 296 objects where 168 are categorized as simple-shaped and 128 as complex-shaped. Example of simple-shaped objects include boxes, cylinders while that of complex-shaped objects involve combination of multiple simple shaped and primitive objects.
5) *Object Segmentation dataset* [35][62] consists of 111 scenes categorized into six subsets i.e. Boxes, Stacked Boxes, Occluded Objects, Cylindrical Objects, Mixed Objects and Complex Scenes. The dataset provides a train/test split of 45 and 66 point clouds respectively. We used the test set as provided by the authors with the dataset.

Baselines. We report results with two baselines. The first baseline involves traditional keypoint based 3D descriptors while the second is am extended deep architecture based on VoxNet.

1) *3D keypoint based descriptors.* We generate supervoxels from the *Universal Training Dataset* and extract various descriptors for points in each supervoxel. The cumulative feature descriptor of the supervoxel is calculated by using average pooling and max pooling of descriptors. The positive and negative pairs of ground-truth data is prepared using the same method as described earlier. We then utilize DBSCAN to find cluster of supervoxels which we mark as discovered objects. The evaluated descriptors are Signatures of Histograms of Orientations (SHOT) [63][64], Fast Point Feature Histograms [65], Rotational Projection Statistics (RoPS)[66] and Spin Image (SI) [67] for there superior performance on various metrics for 3D computer vision tasks [68]. These make for very competitive baselines with traditional hand crafted 3D descriptors.
2) *Deep VoxNet Classification (DVC) Network*. We remove the last layer of the original VoxNet and append two fully connected layers followed by a multiclass classification layer. This method basically classifies each of the training supervoxel into one of the training classes. At test time we pass each of the supervoxel through the network and keep the last fully connected layer responses as the features of the supervoxels. These supervoxel features are then clustered using DBSCAN clustering. This baseline allows us to evaluate the advantages obtained, if any, by using the Siamese proposed method.

Comparison and evaluation setup. On NYU Depth v2 dataset, we compare our results against Stein et al. [41], Silberman et al. [60], and Gupta et al. [42]. The first one is an unsupervised segmentation technique while the latter two involve prior training. The technique of Gupta et al. is related to semantic segmentation in 3D, and is compared here for its good performance on NYU Depth v2 and to evaluate if such techniques can be used for the task of object discovery, i.e. when the training and testing set contain mutually exclusive object categories. For Stein et al. [41] we use the implementation available in Point Cloud Library [69] while for others we use the publicly available code provided by the respective authors. Accuracy is computed by counting the number of objects where the discovered clusters have greater than 80% point-to-point overlap with the ground truth for classes not in the *Universal Training Dataset*. The over-segmentation (F_{os}) and under-segmentation (F_{us}) rates are computed using the method defined in [35] and are given by Eq. 3 where n_{true}, n_{false} and n_{all} are number of correctly classified, incorrectly classified and total number of object points in a point cloud.

$$F_{os} = 1 - \frac{n_{true}}{n_{all}},\ F_{us} = \frac{n_{false}}{n_{all}} \tag{3}$$

We store a mapping of ground-truth labels of points in point cloud to that of supervoxel and further in the voxelized cloud. We use these mappings once the supervoxels have been clustered to obtain n_{true} and n_{false}.

On Object Discovery Dataset, the evaluation metrics used are as defined in [61], which are, over-segmentation (r_{os}), under-segmentation (r_{us}), good-segmentation (r_{gs}) and mis-segmentation (r_{ms}) rates to evaluate various methods on this dataset. The authors compute these metrics by using

labelled segments of object parts. The authors do not provide ground truth segment labelling used to compute these metrics, therefore, we consider each object as a segment for comparison on this dataset, which is a more stringent requirement since the discovered segments need to correspond to entire object instead of object parts. The good-segmentation rate is computed by finding the largest point to point overlap among ground truth and detected objects, while any clusters partially overlapping with the ground truth objects, contribute to the over-segmentation rate. The points in ground-truth object which do not belong to any discovered cluster are considered as mis-segmentations. For remaining discovered clusters, the overlap is calculated with ground-truth clusters, among those the largest overlap cluster is ignored (as they have already been considered), while the rest contribute to mis-segmentation rate.

Training details. We train the VoxNet using the implementation corresponding to [27] on ModelNet10 and ModelNet40 voxelized datasets. We then use this pre-trained VoxNet model with the proposed Siamese architecture.

Discriminative Metric Learning is performed on top of VoxNet as follows. We use the *Universal Training Dataset* described above for training the Siamese network using discriminative metric learning. We use the default parameters specified in the implementation of VoxNet for training the 3D-CNN with ModelNet. The supervoxels are extracted with seed resolutions of 0.05m, 0.10m, 0.15m and 0.2m and subject to the constraint that the % of ground-truth object points in the supervoxel should be greater than or equal to a threshold β. We use the value of $\beta = 0.8$. The value allows for soft-assignment of supervoxels as positive pairs and was chosen based on empirical evaluation described later. The positive set of supervoxels is constructed by pairing all supervoxels on the same object for every object in the training data. The negative set consists of two types of supervoxel pairs. First are the supervoxels belonging to different objects, for this we only use supervoxels around the center of the objects to avoid confusion around boundaries, and second are the pairs of neighbouring supervoxels at object boundaries. We also add pairs of supervoxels belonging to the background classes to the negative set (wall, floor etc.), considering them to contribute to hard negatives as they potentially come from very similar objects.

Testing details. The test set consists of remaining 40% classes of the NYU Depth v2 dataset. The results are reported on classes in test set only even if some of the test scenes may contain classes from training set, this may occur since NYU Depth v2 consists of real world complex scenes and having a clean separation of images with train only and those with test only objects was found to be infeasible.

B. Quantitative Results

The results for proposed VoxNet with Discriminative Metric Learning (VDML) along with the baselines described above is shown in Fig. 5. VDML outperforms other methods, achieving 55.4% and 69.4% test set accuracy on ModelNet10 (VDML10) and ModelNet40 (VDML40) trained 3D-CNN respectively. It is closely followed by DVC with a test set accuracy of 54.7% and 67.3% on ModelNet10 (DVC10) and ModelNet40 (DVC40) respectively. Even training VoxNet with ModelNet10, we observe that the accuracy is nearly same as the DVC baseline while being significantly higher than other baseline methods where the mean performance of descriptors with average pooling is 21.3% while that of max-pooling is 26.6%. Another advantage of VDML over DVC is lower training time. During our experiments, it took 60% more time to train DVC as compared to VDML. This can be explained by the presence of additional fully connected layer in DVC. Compared to Stein et al. , VDML40 achieves 14.9% higher accuracy while method of Silberman et al. and Gupta et al. lag behind by 20.1% and 23.7% respectively. This shows that VDML although following supervised approach is able to generalize better on unseen objects in complex scenes.

Fig. 5 shows over-segmentation (F_{os}) and under-segmentation (F_{us}). The F_{os} for VDML10 and VDML40 are 3.1% and 1.2% respectively. The corresponding values for DVC10 and DVC40 are 3.7% and 1.6% respectively. Comparitively, DVC has higher under-segmentation (5.1% for DVC10 and 2.8% for DVC40) than VDML (3.4% for VDML10 and 1.9% for VDML40). The low F_{os} and F_{us} demonstrate that VDML results in lower cross-overs on object boundaries in segmented regions. The other baselines except DVC perform poorly while methods of Stein et al. , Silberman et al. and and Gupta et al. perform average with F_{os} being 6.1%, 8.2% and 7.4% respectively. An interesting point to note is the significant reduction in F_{os} and F_{us} scores when compared on training on ModelNet10 and ModelNet40 respectively. It shows that when the methods are provided with larger number of models, the resulting features are able to better characterize and learn object boundaries in the embedding space. This underlines the need for more training data in the 3D domain.

The results on Object Discovery Dataset are shown in Table I. As can be seen in the results, the proposed method VDML achieves a good-segmentation rate (r_{os}) of 98.1% on all objects which is a higher by 5.2% when compared to best performing technique (CPC) in [61]. It would be interesting to note that we obtain significantly reduced over-segmentation ($r_{os} = 3.2$) and under-segmentation ($r_{us} = 1.7$) rates on all objects demonstrating that the discovered objects are closer to the complete objects. The mis-segmentation rate, though, is reduced primarily due to background clutter in a few scenes which are not present as ground-truth labels in the training set.

The results on Object Segmentation Dataset are shown in Table II. We achieve a test set accuracy of 99.24% which is 0.83% higher than SVM_{nnb} and 11.2% than SVM_{nb} of [35]. Interestingly, we achieve 100% accuracy on point clouds with geometrically simpler Boxes and Stacked Boxes objects, while also having very low 0% over-segmentation and 0.1% under-segmentation rates. These values become more significant when compared to SVM_{nnb} where higher accuracy is accompanied by significantly large under segmentation rates i.e. 17.2% and 28.2% as compared to VDML40 with 0.1%

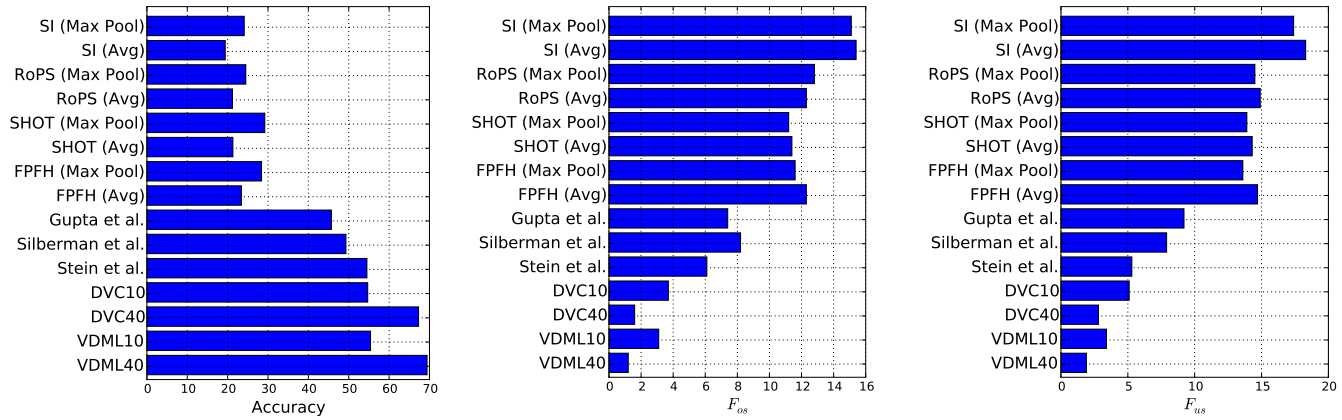

Fig. 5: Accuracies, F_{os} and F_{us} scores of the different methods on NYU Depth v2 dataset. Lower is better for F_{os} and F_{us} scores.

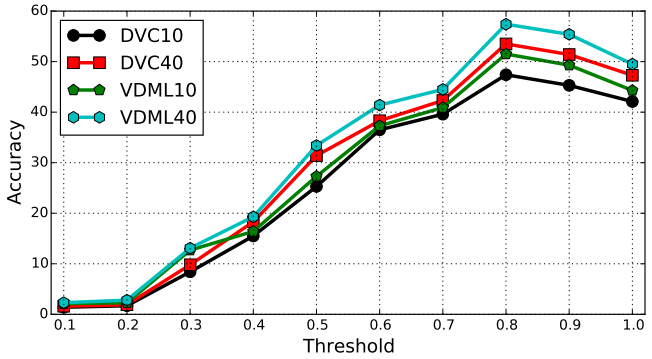

Fig. 4: Accuracy vs. threshold of supervoxel overlap with ground-truth for training on NYU Depth v2 dataset.

and 4.4% on boxes and stacked boxes respectively. The F_{us} rate on complex objects for VDML40 is 4.24% while that for SVM_{nnb} and SVM_{nb} are 146% and 8.0% respectively. SVM_{nnb} witnesses large under-segmentation rate as it is more sensitive to connection between neighbouring segments. On the other hand, VDML is trained on multiple seed resolutions of supervoxels allowing it to learn more intricate connectivity patterns among supervoxel boundaries and hence prevent it from connecting entire objects if a few segments (supervoxels) are incorrectly assigned to same object.

Impact of ground-truth overlap threshold. Fig. 4 shows the variation in accuracy vs. threshold (β), for percentage of points that overlap with ground-truth, to consider it as a positive supervoxel for training. The graph was obtained by splitting the train dataset into train and validation sets with a ratio of 80% to 20% respectively and performing 5-fold cross-validation. It can be observed that the accuracy reaches a peak at 0.8 for both DVC and VDML trained on either of ModelNet10 or ModelNet40. An explanation of better performance of soft supervoxel assignment is because the generated supervoxels do not always lie perfectly within object boundaries. A hard assignment would therefore result in rejection of many supervoxels near the object boundaries. Since our method maximizes the distance between supervoxels on boundaries and neighbouring supervoxels (on background or other objects), a hard assignment has a negative impact on the overall performance which is also validated by a drop of 7.2% and 7.9% (hard-assignment) in accuracy as compared to soft-assignment (β = 0.8) on validation set with VDML10 and VDML40 respectively. Despite having lower validation set accuracy (and subsequently on test set also) DVC on the other hand, witnesses a lower drop of 5.3% (DVC10) and 6.2% (DVC40) as compared to VDML demonstrating that it is relatively less prone supervoxel assignment criteria cf. the proposed VDML.

C. Qualitative Results

Fig. 6 shows visual comparison of result from [35] on Object Segmentation Dataset while Fig. 7 compares to the method from [41]. In Fig. 6 (b), the technique of [35] tends to show under-segmentation and discovers two distinct objects as a single object (grey). In comparison, VDML is able to clearly distinguish among those objects. In Fig. 7, we compare against a scene with numerous occluded objects. Fig. 7(b) shows the results provided by the authors [41]. On comparing Fig. 7(b) and 7(c), it can be seen that VDML results in lesser number of points classified as clutter (shown in black) while providing smoother boundaries which would be helpful for applications such as robot grasping. In both the cases, the can at the front suffers over-segmentation. In the first case, there are two visually arbitrary segments while in the case of the proposed VDML those segments correspond to the body of the can and top of the can. Moreover, the segments from the proposed VDML are closer to the ground-truth as less number of points are marked as clutter.

Fig. 8 and 9 show results of our segmentation on Object Discovery Dataset. There are two important observations to make here. First, as above, the boundaries are closer to the ground-truth. This shows that VDML has learned to distinguish between object and non-object boundaries since the training procedure explicitly focussed on supervoxel pairs around the object boundaries. Secondly, we observe over-segmentation in the object towards the left in Fig 9–the object

T	**Simple-shaped**					**Complex-shaped**					**All Objects**				
N	168					128					296				
S	*APS**	*SPS**	*SPC**	*CPC**	*VDML40*	*APS**	*SPS**	*SPC**	*CPC**	*VDML40*	*APS**	*SPS**	*SPC**	*CPC**	*VDML40*
r_{os}	92.3	43.9	6.6	4.3	2.1	93.1	51.3	22.3	2.8	1.7	92.7	48.8	14.3	4.1	3.2
r_{us}	2.4	1.6	1.9	3.7	2.9	1.3	0.9	0.8	0.8	0.5	1.7	1.4	1.6	2.6	1.7
r_{gs}	6.8	55.2	92.5	92.3	97.8	6.3	48.1	77.1	94.5	98.3	6.6	50.5	84.9	92.9	98.1
r_{ms}	0.8	0.8	0.8	3.4	0.4	0.6	0.6	0.6	2.7	1.6	0.7	0.7	0.7	2.9	1.1

TABLE I: Comparison on Object Discovery Dataset.
*Values from [61], T = Type of Dataset, N = number of objects, S= Segmentation Technique and rates $r_{os,us,gs,ms}$ in %

Technique	SVMnb			SVMnnb			VDML40		
Parameter	**Acc.**	**Fos**	**Fus**	**Acc.**	**Fos**	**Fus**	**Acc.**	**Fos**	**Fus**
Boxes	88.55	1.8	0.2	98.19	0.2	17.2	100.0	0.0	0.1
Stacked Boxes	89.15	1.3	7.1	98.99	0.0	28.2	100.0	0.0	4.4
Occluded Objects	87.93	16.6	0.1	99.23	0.0	0.2	99.96	0.08	0.1
Cylindrical Objects	91.66	2.6	0.3	96.77	2.6	3.5	99.3	0.7	0.9
Mixed Objects	91.04	1.9	19.7	94.97	1.3	1.3	98.5	0.6	1.7
Complex Objects	84.61	7.0	8.0	98.97	5.4	146	97.7	2.3	4.2
Total	87.72	4.5	7.9	98.41	2.7	69.5	**99.24**	**0.61**	**1.9**

TABLE II: Comparison with [35] on Object Segmentation Dataset.

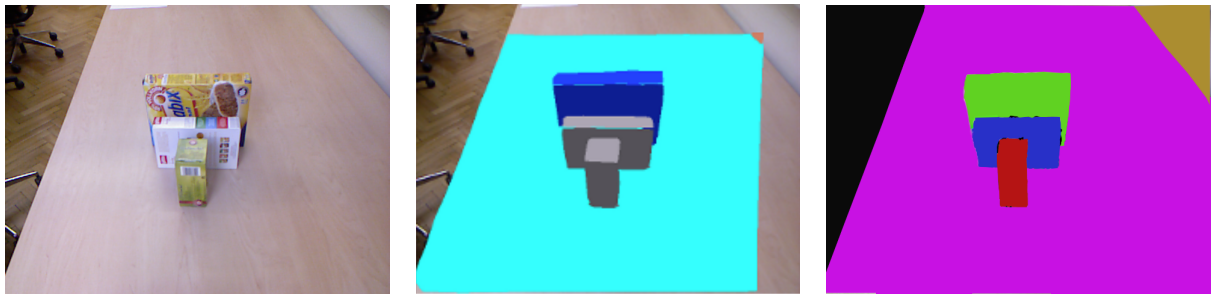

(a) Original Image (b) Result from [35] (c) VDML40

Fig. 6: Comparative result from Object Segmentation Dataset (column 2 has been taken from [35])

(a) Original Image (b) Result from [41] (c) VDML40

Fig. 7: Comparative Result from Object Segmentation Dataset: Column 2 has been taken from [41]. As in [41], for better visualization points after 2m were cropped in the segmented images.

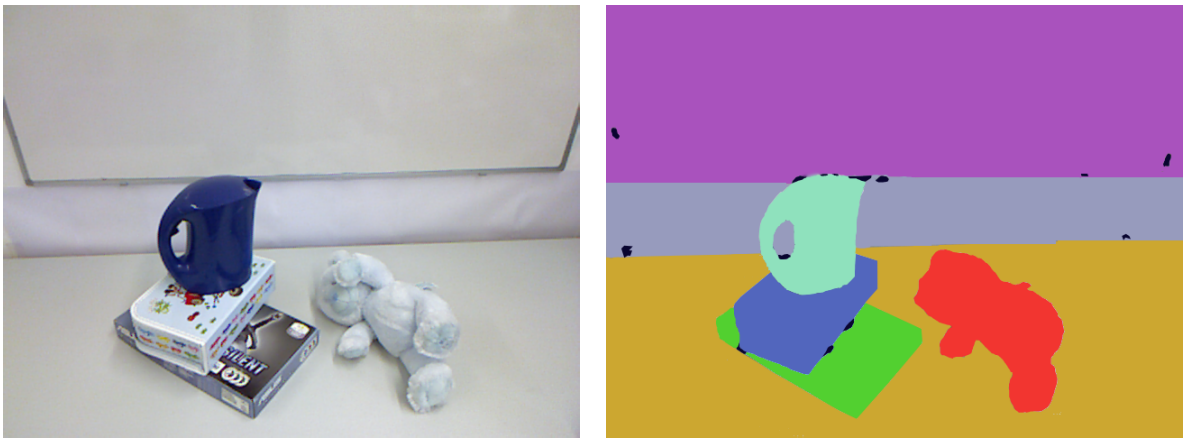

(a) Original Image (b) VDML40

Fig. 8: Qualitative Results on Object Discovery Dataset

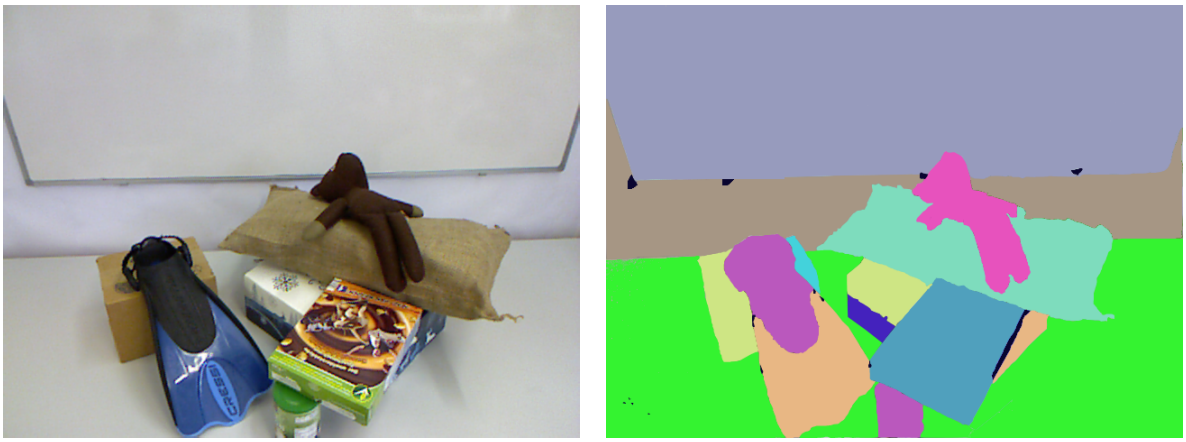

(a) Original Image (b) VDML40

Fig. 9: Qualitative Results on Object Discovery Dataset

(a) Original Image (b) VDML40

Fig. 10: Qualitative Results on NYU Depth v2: Showing results on scene consisting of small objects. (Scene has been cropped for better visualization of the objects)

(a) Original Image (b) VDML40

Fig. 11: Qualitative Results on NYU Depth v2: Showing results on scene consisting of variable sized objects. The images in mirror make it difficult to segment.

gets divided into two segments. Although, we note that the two segments actually belong to visually dissimilar parts of an objects which can easily be considered as distinct objects in the absence of any context.

Fig. 10 and 11 show results on NYU Depth v2 dataset. Since the previous visual images consisted of objects of relatively similar sizes, we now demonstrate the results on scenes with larger variation in sizes of objects. For example, Fig. 11 consists of objects like a soap as well as a human. In the first figure, we see that the objects on the table are extremely cluttered. Despite this, VDML is able to clearly distinguish among the objects on the table (Fig. 10 as well as those in front of the mirror (Fig. 11 except the bottles, where we observe over-segmentation. The wash-basin also gets over-segmented, possibly due to variation in depth and illumination across its surface. The objects in the mirror are similar to the background except the human, which gets over-segmented.

V. Conclusion

We proposed a novel object discovery algorithm which operates in the challenging learning based setting where the objects to be discovered at test time were never seen at train time. We used recent advances in 3D Convolutional Neural Networks and built on it a Siamese deep network for learning non-linear embeddings of supervoxels into a Euclidean space which reflects object based semantics. In the embeddings space the supervoxels from the same objects were constrained to be closer than the supervoxels from different objects. Once these embeddings are learnt they can be used with supervoxels of different novel classes for doing object discovery using efficient clustering in the embedding space. We used a hybrid approach of using CAD models and subsequently RGB-D images for learning non-linear embeddings of supervoxels for discover objects that were never seen before. We provided comprehensive empirical evaluation and comparison with several baselines and existing methods to demonstrate the effectiveness of the technique. We showed that the proposed architecture achieves high accuracy with low over and under segmentation rate making it more reliable in context of competing methods. We also demonstrated qualitatively the improved performance of the proposed method.

References

[1] J. Ren, X. Jiang, J. Yuan, and N. Magnenat-Thalmann, "Sound-event classification using robust texture features for robot hearing," *IEEE Transactions on Multimedia*, 2016.

[2] L. Batrinca, N. Mana, B. Lepri, N. Sebe, and F. Pianesi, "Multimodal personality recognition in collaborative goal-oriented tasks," *IEEE Transactions on Multimedia*, vol. 18, no. 4, pp. 659–673, 2016.

[3] G. B. Keller, T. Bonhoeffer, and M. Hübener, "Sensorimotor mismatch signals in primary visual cortex of the behaving mouse," *Neuron*, vol. 74, no. 5, pp. 809–815, 2012.

[4] A. Wang, J. Lu, J. Cai, T.-J. Cham, and G. Wang, "Large-margin multimodal deep learning for rgb-d object recognition," *IEEE Transactions on Multimedia*, vol. 17, no. 11, pp. 1887–1898, 2015.

[5] J. Tang, L. Jin, Z. Li, and S. Gao, "Rgb-d object recognition via incorporating latent data structure and prior knowledge," *IEEE Transactions on Multimedia*, vol. 17, no. 11, pp. 1899–1908, 2015.

[6] B. Gong, J. Liu, X. Wang, and X. Tang, "Learning semantic signatures for 3d object retrieval," *IEEE Transactions on Multimedia*, vol. 15, no. 2, pp. 369–377, 2013.

[7] H. Saito, S. Baba, and T. Kanade, "Appearance-based virtual view generation from multicamera videos captured in the 3-d room," *IEEE Transactions on Multimedia*, vol. 5, no. 3, pp. 303–316, 2003.

[8] C. Wang, Z. Liu, and S.-C. Chan, "Superpixel-based hand gesture recognition with kinect depth camera," *IEEE transactions on multimedia*, vol. 17, no. 1, pp. 29–39, 2015.

[9] A. Shahroudy, T.-T. Ng, Q. Yang, and G. Wang, "Multimodal multipart learning for action recognition in depth videos," *IEEE transactions on pattern analysis and machine intelligence*, vol. 38, no. 10, pp. 2123–2129, 2016.

[10] Y. Yang, C. Deng, S. Gao, W. Liu, D. Tao, and X. Gao, "Discriminative multi-instance multi-task learning for 3d action recognition," *IEEE Transactions on Multimedia*, 2016.

[11] H. Zhu, J.-B. Weibel, and S. Lu, "Discriminative multi-modal feature fusion for rgbd indoor scene recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 2969–2976.

[12] S. Ma, S. Wang, and W. Gao, "Low complexity adaptive view synthesis optimization in hevc based 3d video coding," *IEEE Transactions on Multimedia*, vol. 16, no. 1, pp. 266–271, 2014.

[13] A. S. Lalos, I. Nikolas, E. Vlachos, and K. Moustakas, "Compressed sensing for efficient encoding of dense 3d meshes using model-based bayesian learning," *IEEE Transactions on Multimedia*, vol. 19, no. 1, pp. 41–53, 2017.

[14] Y. Guo, F. Sohel, M. Bennamoun, J. Wan, and M. Lu, "An accurate and robust range image registration algorithm for 3d object modeling," *ieee transactions on multimedia*, vol. 16, no. 5, pp. 1377–1390, 2014.

[15] J.-Y. Chen, C.-H. Lin, P.-C. Hsu, and C.-H. Chen, "Point cloud encoding for 3d building model retrieval," *IEEE Transactions on Multimedia*, vol. 16, no. 2, pp. 337–345, 2014.

[16] D. Philipona, J. K. O'Regan, and J.-P. Nadal, "Is there something out there? inferring space from sensorimotor dependencies," *Neural computation*, vol. 15, no. 9, pp. 2029–2049, 2003.

[17] S. Gupta, R. Girshick, P. Arbeláez, and J. Malik, "Learning rich features from rgb-d images for object detection and segmentation," in *European Conference on Computer Vision*. Springer, 2014, pp. 345–360.

[18] S. Gupta, P. Arbeláez, R. Girshick, and J. Malik, "Indoor scene understanding with rgb-d images: Bottom-up segmentation, object detection and semantic segmentation," *International Journal of Computer Vision*, vol. 112, no. 2, pp. 133–149, 2015.

[19] T. Tuytelaars, C. H. Lampert, M. B. Blaschko, and W. Buntine, "Unsupervised object discovery: A comparison," *International journal of computer vision*, vol. 88, no. 2, pp. 284–302, 2010.

[20] A. Karpathy, S. Miller, and L. Fei-Fei, "Object discovery in 3d scenes via shape analysis," in *Robotics and Automation (ICRA), 2013 IEEE International Conference on*. IEEE, 2013, pp. 2088–2095.

[21] C. A. Mueller and A. Birk, "Hierarchical graph-based discovery of non-primitive-shaped objects in unstructured environments," in *Robotics and Automation (ICRA), 2016 IEEE International Conference on*. IEEE, 2016, pp. 2263–2270.

[22] M. Firman, "RGBD Datasets: Past, Present and Future," in *CVPR Workshop on Large Scale 3D Data: Acquisition, Modelling and Analysis*, 2016.

[23] Z. Wu, S. Song, A. Khosla, F. Yu, L. Zhang, X. Tang, and J. Xiao, "3d shapenets: A deep representation for volumetric shapes," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 1912–1920.

[24] S. Song and J. Xiao, "Deep Sliding Shapes for amodal 3D object detection in RGB-D images," in *CVPR*, 2016.

[25] X. Chen, K. Kundu, Y. Zhu, A. G. Berneshawi, H. Ma, S. Fidler, and R. Urtasun, "3d object proposals for accurate object class detection," in *Advances in Neural Information Processing Systems*, 2015, pp. 424–432.

[26] A. Eitel, J. T. Springenberg, L. Spinello, M. Riedmiller, and W. Burgard, "Multimodal deep learning for robust rgb-d object recognition," in *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*. IEEE, 2015, pp. 681–687.

[27] D. Maturana and S. Scherer, "Voxnet: A 3d convolutional neural network for real-time object recognition," in *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*. IEEE, 2015, pp. 922–928.

[28] S. Bu, Z. Liu, J. Han, J. Wu, and R. Ji, "Learning high-level feature by deep belief networks for 3-d model retrieval and recognition," *IEEE Transactions on Multimedia*, vol. 16, no. 8, pp. 2154–2167, 2014.

[29] F. Schroff, D. Kalenichenko, and J. Philbin, "Facenet: A unified embedding for face recognition and clustering," in *CVPR*, 2015.

[30] B. Bhattarai, G. Sharma, and F. Jurie, "CP-mtML: Coupled projection multi-task metric learning for large scale face retrieval," in *CVPR*, 2016.

[31] C. Huang, C. Change Loy, and X. Tang, "Unsupervised learning of discriminative attributes and visual representations," in *CVPR*, 2016.

[32] X. Wang and A. Gupta, "Unsupervised learning of visual representations using videos," in *ICCV*, 2015.

[33] G. M. García, E. Potapova, T. Werner, M. Zillich, M. Vincze, and S. Frintrop, "Saliency-based object discovery on rgb-d data with a late-fusion approach," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 1866–1873.

[34] J. Papon, A. Abramov, M. Schoeler, and F. Worgotter, "Voxel cloud connectivity segmentation-supervoxels for point clouds," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2013, pp. 2027–2034.

[35] A. Richtsfeld, T. Mörwald, J. Prankl, M. Zillich, and M. Vincze, "Segmentation of unknown objects in indoor environments," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2012, pp. 4791–4796.

[36] M. Firman, D. Thomas, S. Julier, and A. Sugimoto, "Learning to discover objects in rgb-d images using correlation clustering," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2013, pp. 1107–1112.

[37] E. Herbst, X. Ren, and D. Fox, "Rgb-d object discovery via multi-scene analysis," in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2011, pp. 4850–4856.

[38] E. Herbst, P. Henry, X. Ren, and D. Fox, "Toward object discovery and modeling via 3-d scene comparison," in *Robotics and Automation (ICRA), 2011 IEEE International Conference on*. IEEE, 2011, pp. 2623–2629.

[39] J. Shin, R. Triebel, and R. Siegwart, "Unsupervised discovery of repetitive objects," in *ICRA*, 2010, pp. 5041–5046.

[40] J. Bao, Y. Jia, Y. Cheng, and N. Xi, "Saliency-guided detection of unknown objects in rgb-d indoor scenes," *Sensors*, vol. 15, no. 9, pp. 21 054–21 074, 2015.

[41] S. Christoph Stein, M. Schoeler, J. Papon, and F. Worgotter, "Object partitioning using local convexity," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 304–311.

[42] S. Gupta, P. Arbelaez, and J. Malik, "Perceptual organization and recognition of indoor scenes from rgb-d images," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2013, pp. 564–571.

[43] U. Asif, M. Bennamoun, and F. Sohel, "Unsupervised segmentation of unknown objects in complex environments," *Autonomous Robots*, vol. 40, no. 5, pp. 805–829, 2016.

[44] Q. Zhang, X. Song, X. Shao, H. Zhao, and R. Shibasaki, "Unsupervised 3d category discovery and point labeling from a large urban environment," in *Robotics and Automation (ICRA), 2013 IEEE International Conference on*. IEEE, 2013, pp. 2685–2692.

[45] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Süsstrunk, "Slic superpixels compared to state-of-the-art superpixel methods," *IEEE transactions on pattern analysis and machine intelligence*, vol. 34, no. 11, pp. 2274–2282, 2012.

[46] K. Xu, H. Huang, Y. Shi, H. Li, P. Long, J. Caichen, W. Sun, and B. Chen, "Autoscanning for coupled scene reconstruction and proactive object analysis," *ACM Transactions on Graphics (TOG)*, vol. 34, no. 6, p. 177, 2015.

[47] F. Ponjou Tasse, J. Kosinka, and N. Dodgson, "Cluster-based point set saliency," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 163–171.

[48] Y. Shi, P. Long, K. Xu, H. Huang, and Y. Xiong, "Data-driven contextual modeling for 3d scene understanding," *Computers & Graphics*, vol. 55, pp. 55–67, 2016.

[49] A. P. Moore, S. J. Prince, J. Warrell, U. Mohammed, and G. Jones, "Superpixel lattices," in *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*. IEEE, 2008, pp. 1–8.

[50] M. Van den Bergh, X. Boix, G. Roig, B. de Capitani, and L. Van Gool, "Seeds: Superpixels extracted via energy-driven sampling," in *European conference on computer vision*. Springer, 2012, pp. 13–26.

[51] O. Veksler, Y. Boykov, and P. Mehrani, "Superpixels and supervoxels in an energy optimization framework," in *European conference on Computer vision*. Springer, 2010, pp. 211–224.

[52] D. Weikersdorfer, D. Gossow, and M. Beetz, "Depth-adaptive superpixels," in *Pattern Recognition (ICPR), 2012 21st International Conference on*. IEEE, 2012, pp. 2087–2090.

[53] D. Maturana and S. Scherer, "3d convolutional neural networks for landing zone detection from lidar," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 3471–3478.

[54] A. Garcia-Garcia, F. Gomez-Donoso, J. Garcia-Rodriguez, S. Orts-Escolano, M. Cazorla, and J. Azorin-Lopez, "Pointnet: A 3d convolutional neural network for real-time object class recognition," in *Neural Networks (IJCNN), 2016 International Joint Conference on*. IEEE, 2016, pp. 1578–1584.

[55] J. Huang and S. You, "Point cloud labeling using 3d convolutional neural network," in *Proc. of the International Conf. on Pattern Recognition (ICPR)*, vol. 2, 2016.

[56] C. R. Qi, H. Su, M. Niessner, A. Dai, M. Yan, and L. J. Guibas, "Volumetric and multi-view cnns for object classification on 3d data," in *CVPR*, 2016.

[57] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf, "Deepface: Closing the gap to human-level performance in face verification," in *CVPR*, 2014.

[58] G. Sharma and B. Schiele, "Scalable nonlinear embeddings for semantic category-based image retrieval," in *ICCV*, 2015.

[59] M. Ester, H.-P. Kriegel, J. Sander, X. Xu *et al.*, "A density-based algorithm for discovering clusters in large spatial databases with noise." in *Kdd*, vol. 96, no. 34, 1996, pp. 226–231.

[60] P. K. Nathan Silberman, Derek Hoiem and R. Fergus, "Indoor segmentation and support inference from rgbd images," in *ECCV*, 2012.

[61] C. A. Mueller and A. Birk, "Hierarchical Graph-Based Discovery of Non-Primitive-Shaped Objects in Unstructured Environments," in *IEEE International Conference on Robotics and Automation (ICRA)*, May 2016.

[62] "Andreas richtsfeld. the object segmentation database (osd)," http://www.acin.tuwien.ac.at/?id=289, accessed: 2017-01-06.
[63] F. Tombari, S. Salti, and L. Di Stefano, "Unique signatures of histograms for local surface description," in *European conference on computer vision*. Springer, 2010, pp. 356–369.
[64] S. Salti, F. Tombari, and L. Di Stefano, "Shot: unique signatures of histograms for surface and texture description," *Computer Vision and Image Understanding*, vol. 125, pp. 251–264, 2014.
[65] R. B. Rusu, N. Blodow, and M. Beetz, "Fast point feature histograms (fpfh) for 3d registration," in *Robotics and Automation, 2009. ICRA'09. IEEE International Conference on*. IEEE, 2009, pp. 3212–3217.
[66] Y. Guo, F. Sohel, M. Bennamoun, M. Lu, and J. Wan, "Rotational projection statistics for 3d local surface description and object recognition," *International journal of computer vision*, vol. 105, no. 1, pp. 63–86, 2013.
[67] A. E. Johnson and M. Hebert, "Using spin images for efficient object recognition in cluttered 3d scenes," *IEEE Transactions on pattern analysis and machine intelligence*, vol. 21, no. 5, pp. 433–449, 1999.
[68] Y. Guo, M. Bennamoun, F. Sohel, M. Lu, J. Wan, and N. M. Kwok, "A comprehensive performance evaluation of 3d local feature descriptors," *International Journal of Computer Vision*, vol. 116, no. 1, pp. 66–89, 2016.
[69] R. B. Rusu and S. Cousins, "3d is here: Point cloud library (pcl)," in *Robotics and Automation (ICRA), 2011 IEEE International Conference on*. IEEE, 2011, pp. 1–4.